

O

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

**Each of DIDS February 1991 and DIDS October 1991 invalidate the indicated claims under 35 U.S.C. § 102(b)
and 35 U.S.C. § 103^{*}**

All text citations for the “printed publication” entitled “DIDS February 1991” are taken from: Steven Snapp et al., “Intrusion Detection Systems (IDS): A Survey of Existing Systems and A Proposed Distributed IDS Architecture” (February 1991) [SYM_P_0069280- SYM_P_0069297].

All text citations for the “printed publication” entitled “DIDS October 1991” are taken from: S.R. Snapp, J. Brentano, G.V. Dias, L.T. Heberlein, C. Ho, K.N. Levitt, B. Mukherjee, (with S.E. Smaha, T. Grance, D.M. Teal, D.L. Mansur), “DIDS -- Motivation, Architecture, and an Early Prototype” Proc. 14th National Computer Security Conference, Washington, DC, Oct. 1991, pp. 167-176 [SYM_P_0077175- SYM_P_0077185]. SRI has admitted this paper was published before November 9, 1997. See SRI's Responses to Symantec's Third Set of RFAs, #19.

The text included herein are merely representative samples of the disclosure in the asserted reference. I reserve the right to supplement these disclosures.

Similar disclosures and additional related information are contained in the following additional references:

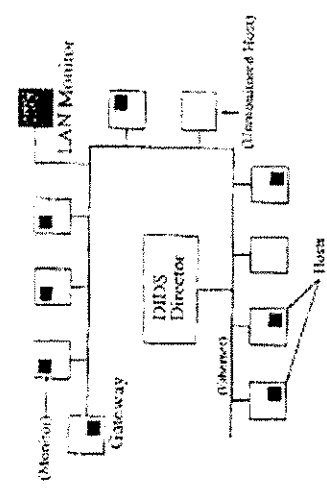
- S.R. Snapp, J. Brentano, G.V. Dias, T.L. Goan, L.T. Heberlein, C.L. Ho, K.N. Levitt, B. Mukherjee, (with T. Grance D.L. Mansur, K.L. Pon, S.E. Smaha), “A System for Distributed Intrusion Detection,” COMPCON Spring 91, Digest of Papers, San Francisco, CA, 25 Feb.-1 March 1991, pp. 170-176 [SYM_P_0069210- SYM_P_0069216].
- J. Brentano, S.R. Snapp, G.V. Dias, T.L. Goan, L.T. Heberlein, C.L. Ho, K.N. Levitt, B. Mukherjee, (with S.E. Smaha), “An Architecture for a Distributed Intrusion Detection System,” Proc. of the 14th Department of Energy Computer Security Group Conference, May 1991, pp.(17)25-(17)45 [SYM_P_0069217- SYM_P_0069238],

^{*} 103 references are identified under the heading “103.”

Distributed Intrusion Detection System “DIDS February 1991 and DIDS October 1991”

- Steve Snapp, “Signature analysis and communication issues in a distributed intrusion detection system,” M.S. Thesis, Division of Computer Science, University of California, Davis, August 1991 [SYM_P_0069163- SYM_P_0069209].
- S.R. Snapp, J. Brentano, G.V. Dias, L.T. Heberlein, C. Ho, K.N. Levitt, B. Mukherjee, (with S.E. Smaha, T. Grance, D.M. Teal, D.L. Mansur), “DIDS -- Motivation, Architecture, and an Early Prototype,” Proc. 14th National Computer Security Conference, Washington, DC, Oct. 1991, pp. 167-176 [SYM_P_0077175- SYM_P_0077185].
- Steven R. Snapp, Stephen E. Smaha, Daniel M. Teal, Tim Grance, “The DIDS (Distributed Intrusion Detection System) Prototype,” Proceedings of the Summer 1992 USENIX Conference, June 8-12, 1992 [SYM_P_0501723- SYM_P_0501736].
- B. Mukherjee, L.T. Heberlein, K.N. Levitt, “Network Intrusion Detection,” IEEE Network, May-June 1994, Vol. 8, No. 3, pp. 26-41 [SYM_P_0069263- SYM_P_0069279].
- Terrance Lee Goan Jr., “Towards a Dynamic System for Accountability and Intrusion Detection in a Networked Environment,” M.S. Thesis, Division of Computer Science, University of California, Davis, 1992 [SYM_P_0598736-95].
- Justin Edgar Doak, “Intrusion Detection: the Application of Feature Selection, a Comparison of Algorithms, and the Application of a Wide Area Network Analyzer,” M.S. Thesis, Division of Computer Science, University of California, Davis, 1992 [SYM_P_0598736-95]

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
1	A method of network surveillance, comprising:	<p>“The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“We believe that DIDS will be able to detect the same kind of single host intrusions that are flagged by other intrusion detection systems, such as IDES [6], Wisdom & Sense [15], and MIPAS [10]. DIDS should also be able to (1) detect attacks on the network itself; (2) detect attacks involving multiple hosts; (3) track tagged objects, including users and sensitive files, as they move around the network; (4) detect, via erroneous or misleading reports, situations where a host might be taken over by an attacker; and (5) monitor the activity of any networked system that doesn't have a host monitor, yet generates LAN activity, such as a PC.” (15) [SYM_P_0069294]</p>	<p>“We are designing and implementing a prototype Distributed Intrusion Detection System (DIDS) that combines distributed monitoring and data reduction (through individual host and LAN monitors) with centralized data analysis (through the DIDS director) to monitor a heterogeneous network of computers. This approach is unique among current IDS's.” (167) [SYM_P_0077175]</p>  <p style="text-align: center;">Fig. 1. DIDS Target Environment</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>[SYM_P_0069297]</p> <p style="text-align: center;">Fig. 4 Architecture of the Distributed Intrusion Detection System</p>	<p>[SYM_P_0077184]</p> <p style="text-align: center;">Fig. 3 Communications Architecture</p> <p style="text-align: center;">Fig. 5 Host Monitor Structure</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
	receiving network packets handled by a network entity;	<p>“Previous work on intrusion-detection systems were performed on stand-alone hosts and on a broadcast local area network (LAN) environment. The focus of our present research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well. The generalized distributed environment is heterogeneous, i.e., the network nodes can be hosts or servers from different vendors, or some of them could be LAN managers, like our previous work, a network security monitor (NSM), as well. The proposed architecture for this distributed intrusion-detection system consists of the following components: a host manager (viz. a monitoring process or collection of processes running in background) in each host; a LAN manager for monitoring each LAN in the system; and a central manager which is placed at a single secure location and which receives reports from various host and LAN managers to process these reports, correlate them, and detect intrusions.” (1) [SYM_P_0069280]</p> <p>“The LAN monitor sees every packet on its segment of the LAN.” (13) [SYM_P_0069292]</p> <p>“The <i>Network Security Monitor</i> (NSM) is different from the</p>	<p>[SYM_P_0077184]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.” (171) [SYM_P_0077179]</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>intrusion detection systems discussed in Section 2 in that it does not analyze audit trails to detect intrusive behavior. The NSM, as the name implies, analyzes the traffic on a broadcast local area network to detect intrusive behavior. The reasons for this departure from the standard intrusion detection methods are outlined as follows.</p> <p>“First, although most IDSs are designed with the goal of supporting a number of different operating system platforms, all present audit-trail-based IDSs have only been used on a single operating system at any one time. These systems are usually designed to transform an audit log into a proprietary format used by the IDS [6, 10, 11]. In theory, audit logs from different operating systems need only to be transformed into this proprietary form for the IDS to perform its analysis. However, no results of an IDS successfully supporting multiple operating systems have been reported.</p> <p>“On the other hand, standard network protocols exist (e.g., TCP/IP and UDP/IP) which most major operating systems support and use. By using these network standards, the NSM can monitor a heterogeneous set of hosts and operating systems simultaneously.</p> <p>“Second, audit trails are often not available in a timely fashion. Some IDSs are designed to perform their analysis on a separate</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>host, so the audit logs must be transferred from the source host to the second host monitor [11]. Furthermore, the operating system can often delay the writing of audit logs by several minutes [15]. The broadcast nature of local area networks, however, gives the NSM instant access to all data as soon as this data is transmitted on the network. It is then possible to immediately start the attack detection process.</p> <p>“Third, the audit trails are often vulnerable. In some past incidents, the intruders have turned off audit daemons or modified the audit trail. This action can either prevent the detection of the intrusion, or it can remove the capability to perform accountability (who turned off the audit daemons?) and damage control (what was seen, modified, or destroyed?) The NSM, on the other hand, passively listens to the network, and is therefore logically protected from subversion. Since the NSM is invisible to the intruder, it cannot be turned off (assuming it is physically secured), and the data it collects cannot be modified.</p> <p>“Fourth, the collection of audit trails degrades the performance of a machine being monitored. Unless audit trails are being used for accounting purposes, system administrators often turn off auditing. If analysis of these audit logs is also to be performed on the host, added degradation will occur. If the audit logs are transferred across a network or communication channel to a separate host for analysis, the loss of network bandwidth, as well</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
	building at least one long-term and at least one short-term statistical profile from at least one measure of the network	<p>as the loss of timeliness of the data will occur. In many environments, the degradation of monitored hosts or the loss of network bandwidth may discourage administrators from using such an IDS. The alternative, viz. the NSM architecture, does not degrade the performance of the hosts being monitored. The monitored hosts are not aware of the NSM, so the effectiveness of the NSM is not dependent on the system administrator's configuration of the monitored hosts.</p> <p>“And, finally, many of the more seriously documented cases of computer intrusions have utilized a network at some point during the intrusion, i.e., the intruder was physically separated from the target. With the continued proliferation of networks and interconnectivity, the use of networks in attacks will only increase. Furthermore, the network itself, being an important component of a computing environment, can be the object of an attack. The NSM can take advantage of the increase of network usage to protect the hosts attached to the network. It can monitor attacks launched against the network itself, an attack that host based audit trail analyzers would probably miss.” (9-10) [SYM_P_0069288- SYM_P_0069289]</p>	<p>“The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
	packets	<p>items as login times, CPU usage, favorite editor and compiler, disk usage, number of printed pages per session, session length, error rate, etc. (T.F. Lunt et al., [7] presents a comprehensive list of possible measures.) The IDS will then use these profiles to compare current user activity with past user activity. Whenever a current user's activity pattern falls outside certain pre-defined thresholds, the behavior is considered anomalous. Legitimate behavior that is flagged as intrusive is defined to be a <i>false alarm</i>. A major problem with the statistical method is determining exactly what activities and statistical measures provide the highest detection rate and lowest false alarm rate for a particular system. Those statistics that detect an attack on a computer system may differ from system to system depending on the system and its environment; so the measures must be tailored for each particular system. It may also be the case that a particular activity may not be threatening by itself, but when aggregated with other activities, it may constitute an attack. These statistical profiles must be adaptive, i.e., they must be updated regularly, since users may be constantly changing their behavior.” (2) [SYM_P_0069281]</p> <p>“The traffic on the network is analyzed by a simple expert system. The types of inputs to the expert system are described below.</p> <p>“The current traffic cast into the ICEM vectors as discussed in</p>	<p>profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>the previous subsection is the first type of input. Currently, only the connection vectors and the host vectors are used. The components for these vectors are presented in Tables I and II.</p> <p>“The profiles of expected traffic behavior are the second type input. The profiles consist of expected data paths (viz. which systems are expected to establish communication paths to which other systems, and by which service?) and service profiles (viz. what is a typical <i>telnet</i>, <i>mail</i>, <i>finger</i>, etc., expected to look like?). Combining profiles and current network traffic gives the NSM the ability to detect anomalous behavior on the network.</p> <p>“The knowledge about capabilities of each of the network services is the third type of input (e.g., <i>telnet</i> provides the user with more capability than <i>ftp</i> does).</p> <p>“The level of authentication required for each of the services is the fourth type of input (e.g., <i>finger</i> requires no authentication, <i>mail</i> requests authentication but does not verify it, and <i>telnet</i> requires verified authentication).</p> <p>“The level of security for each of the machines is the fifth type of input. This can be based on the NCSC rating of machines, history of past abuses on the different machines, the rating received after running system evaluation software such as SPI or COPS, or simply which machines the security officer has some</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>control over and which machines the security officer has no control over (e.g., a host from outside the local area network).</p> <p>“And the signatures of past attacks is the sixth type of input. Examples include seeing the vertical bar symbol (i.e., , a Unix “pipe” symbol) in the receiver address for <i>mail</i>, or <i>finger</i> connections where the initiating host sends more than 512 bytes to the receiving host.</p> <p>“The data from these sources is used to identify the likelihood that a particular connection represents intrusive behavior, or if a host has been compromised. The security state, or suspicion level, of a particular connection is a function of the abnormality of the connection, the security level of the service being used for the connection, the direction of the connection security level, and the matched signatures of attacks in the data stream for that connection.</p> <p>“The abnormality of a connection is based on the probability of that particular connection occurring and the behavior of the connection itself. If a connection from host A to host B by service C is rare, then the abnormality of that connection is high. Furthermore, if the profile of that connection compared to a typical connection by the same type of service is unusual (e.g., the number of packets or bytes is unusually high for a <i>mail</i> connection), the abnormality of that connection is high.</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>“The security level of the service is based on the capabilities of that service and the authentication required by that service. The <i>tfip</i> service, for example, has great capabilities with no authentication, so the security level for <i>tfip</i> is high. The <i>telnet</i> service, on the other hand, also has great capabilities, but it also requires strong authentication. Therefore, the security level for <i>telnet</i> is <i>lower</i> than that of <i>tfip</i>.”</p> <p>“The direction of connection security level is based on the security levels of the two machines involved and which host initiated the connection. If a low security host connects to, or attempts to connect to a high security host, the direction of connection security level of that connection is high. On the other hand, if a high security host connects to an insecure host, the direction of connection security level is low.</p> <p>“The matched strings consists of the vectors Initiator_X and Receiver_X. Thus it is simply a list of counts for the number of times each string being searched for in the data is matched.</p> <p>“The connection vectors are essentially treated as records in a database, and presentation of the information may be made as simple requests into the database. The default presentation format sorts the connection by suspicion level and presents the sorted list from highest suspicion level to the lowest.</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>Presentations can also be made by specifying time windows for connection, connections from a specific host, connections with a particular string matched, etc.</p> <p>“The security state, or suspicion level, of a host is simply the maximum security state of its connection vectors over particular window of time. The host vectors are also treated as records into a database, and they may be presented in a similar fashion as the connection vectors.” (10-11) [SYM_P_0069289- SYM_P_0069290]</p> <p>“4.5. LAN Monitor</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director.</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>It also uses and maintains profiles of network behavior, which are updated periodically. Like the host monitor, the LAN monitor provides an agent for communications with the director. In addition to handling queries of the audit data from the director, this agent gives the director access to a number of network management tools, which are analogous to the host operating system services provided by the host monitor.” (13) [SYM_P_0069292]</p>	
	<p>the at least one measure monitoring data transfers, errors, or network connections;</p>	<p>“The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic.” (13) [SYM_P_0069292]</p> <p>“The NSM models the network and hosts being monitored in the hierarchically-structured Interconnected Computing Environment Model (ICEM). The ICEM is composed of six layers, the lowest being the bit streams on the network, and the highest being a representation for the state of the entire networked system.</p> <p>The bottom-most, or first, layer is the <i>packet layer</i>. This layer accepts as input a <i>bit stream</i> from a broadcast local area network, viz. an Ethernet. The bit stream is divided up into complete Ethernet packets, and a time stamp is attached to the</p>	<p>“Like the host monitor, the LAN monitor consists of a <i>LAN event generator</i> (LEG) and a <i>LAN agent</i>. The LEG is currently a subset of UC Davis’ NSM [3]. Its main responsibility is to observe all of the traffic on its segment of the LAN to monitor host-to-host connections, services used, and volume of traffic. The LAN monitor reports on such network activity as <i>rlogin</i> and <i>telnet</i> connections, the use of security-related services, and changes in network traffic patterns.” (169) [SYM_P_0077177]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.” (171)</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>packet. This <i>time-augmented packet</i> is then passed up to the second layer.</p> <p>The next layer, called the <i>thread layer</i>, accepts as input the time-augmented packets from the packet layer. These packets are then correlated into unidirectional data streams. Each stream consists of the data (with the different layers of protocol headers removed) being transferred from one host to another host by a particular protocol (TCP/IP or UDP/IP), through a unique set (for the particular set of hosts and protocol) of ports. This stream of data, which is called a thread, is mapped into a <i>thread vector</i>. All the thread vectors are passed up to the third layer.</p> <p>The <i>connection layer</i>, which is the third layer, accepts as input the thread vectors generated by the thread layer. Each thread vector is paired, if possible, to another thread vector to represent a bidirectional stream of data (i.e., a host-to-host connection). These pairs of thread vectors are represented by a connection vector generated by the combination of the individual thread vectors. Each connection vector will be analyzed, and a <i>reduced</i> representation, a <i>reduced connection vector</i>, is passed up to the fourth layer.</p> <p>Layer 4 is the <i>host layer</i> which accepts as input the reduced connection vectors generated by the connection layer. The connection vectors are used to build <i>host vectors</i>. Each host</p>	<p>[SYM_P_0077179]</p> <p>“[3] L.T. Heberlein, G. Dias, K. Levitt, B. Mukherjee, J. Wood, and D. Wolber, ‘A Network Security Monitor,’ <i>Proc. 1990 Symposium on Research in Security and Privacy</i>, pp. 296-2304, Oakland, CA, May 1990.” [SYM_P_0077183]</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>vector represents the network activities of a single host. These host vectors are passed up to the fifth layer.</p> <p>The <i>connected network</i> layer is the next layer in the ICEM hierarchy. It accepts as input the host vectors generated by the host layer. The host vectors are transformed into a graph G by treating the Data_path_tuples of the host vectors as an adjacency list. If G (host1,host2,serv1) is not empty, then there is a connection, or path, from host1 to host2 by service serv1. The value for location G(host1,host2,serv1) is non empty if the host vector for host1 has (host2,serv1) in its Data_path_tuples. This layer can build the connected sub-graphs of G, called a <i>connected network vector</i>, and compare these sub-graphs against historical connected sub-graphs. This layer can also accept questions from the user about the graph. For example, the user may ask if there is some path between two hosts through any number of intermediate hosts — by a specific service. This set of connected network vectors is passed up to the sixth and final layer.</p> <p>The top most layer, called the <i>system layer</i>, accepts as input the set of connected network vectors from the connected network layer. The set of connected network vectors are used to build a <i>single system</i> vector representing the behavior of the entire system.” (10) [SYM_P_0069289]</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically. Like the host monitor, the LAN</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
	<p>comparing at least one long-term and at least one short-term statistical profile; and</p>	<p>monitor provides an agent for communications with the director. In addition to handling queries of the audit data from the director, this agent gives the director access to a number of network management tools, which are analogous to the host operating system services provided by the host monitor.” (13) [SYM_P_0069292]</p> <p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The abnormality of a connection is based on the probability of that particular connection occurring and the behavior of the connection itself. If a connection from host A to host B by service C is rare, then the abnormality of that connection is high. Furthermore, if the profile of that connection compared to a typical connection by the same type of service is unusual (e.g., the number of packets or bytes is unusually high for a <i>mail</i> connection), the abnormality of that connection is high.” (11)</p>	<p>“The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>“Like the host monitor, the LAN monitor consists of a <i>LAN event generator</i> (LEG) and a <i>LAN agent</i>. The LEG is currently a subset of UC Davis’ NSM [3]. Its main responsibility is to observe all of the traffic on its segment of the LAN to monitor host-to-host connections, services used, and volume of traffic. The LAN monitor reports on such network activity as <i>rlogin</i> and <i>telnet</i> connections, the use of security-related services, and changes in network traffic patterns.” (169) [SYM_P_0077177]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>[SYM_P_0069290]</p> <p>“The profiles of expected traffic behavior are the second type input. The profiles consist of expected data paths (viz. which systems are expected to establish communication paths to which other systems, and by which service?) and service profiles (viz. what is a typical <i>telnet</i>, <i>mail</i>, <i>finger</i>, etc., expected to look like?)</p> <p>Combining profiles and current network traffic gives the NSM the ability to detect anomalous behavior on the network.” (10)</p> <p>[SYM_P_0069289- SYM_P_0069290]</p>	<p>on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.” (171)</p> <p>[SYM_P_0077179]</p> <p>“[3] L.T. Heberlein, G. Dias, K. Levitt, B. Mukherjee, J. Wood, and D. Wolber, ‘A Network Security Monitor,’ <i>Proc. 1990 Symposium on Research in Security and Privacy</i>, pp. 296-2304, Oakland, CA, May 1990.” [SYM_P_0077183]</p>
determining whether the difference between the short-term statistical profile and the long-term statistical profile indicates suspicious network activity.		<p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The abnormality of a connection is based on the probability of that particular connection occurring and the behavior of the connection itself. If a connection from host A to host B by</p>	<p>“The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>“Like the host monitor, the LAN monitor consists of a <i>LAN event generator</i> (LEG) and a <i>LAN agent</i>. The LEG is currently a subset of UC Davis’ NSM [3]. Its main responsibility is to observe all of the traffic on its segment of the LAN to monitor host-to-host connections, services used, and volume of traffic. The LAN monitor reports on such network activity as <i>rlogin</i> and <i>telnet</i> connections, the use of security-related services, and</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>service C is rare, then the abnormality of that connection is high. Furthermore, if the profile of that connection compared to a typical connection by the same type of service is unusual (e.g., the number of packets or bytes is unusually high for a <i>mail</i> connection), the abnormality of that connection is high.” (11) [SYM_P_0069290]</p> <p>“The profiles of expected traffic behavior are the second type [sic] input. The profiles consist of expected data paths (viz. which systems are expected to establish communication paths to which other systems, and by which service?) and service profiles (viz. what is a typical <i>telnet</i>, <i>mail</i>, <i>finger</i>, etc., expected to look like?) Combining profiles and current network traffic gives the NSM the ability to detect anomalous behavior on the network.” (10-11) [SYM_P_0069289- SYM_P_0069290]</p>	<p>changes in network traffic patterns.” (169) SYM_P_0077177]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.” (171) [SYM_P_0077179]</p> <p>“[3] L.T. Heberlein, G. Dias, K. Levitt, B. Mukherjee, J. Wood, and D. Wolber, ‘A Network Security Monitor,’ <i>Proc. 1990 Symposium on Research in Security and Privacy</i>, pp. 296-2304, Oakland, CA, May 1990.” [SYM_P_0077183]</p>

Distributed Intrusion Detection System "DIDS February 1991 and DIDS October 1991"

Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
338	Claim	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
2	The method of claim 1, wherein the measure monitors data transfers by monitoring network packet data transfer commands	<p><u>103:</u></p> <p>“The knowledge about capabilities of each of the network services is the third type of input (e.g., <i>telnet</i> provides the user with more capability than <i>ftp</i> does).” (11) [SYM_P_0069290]</p> <p>“The security level of the service is based on the capabilities of that service and the authentication required by that service. The <i>ftp</i> service, for example, has great capabilities with no authentication, so the security level for <i>ftp</i> is high. The <i>telnet</i> service, on the other hand, also has great capabilities, but it also requires strong authentication. Therefore, the security level for <i>telnet</i> is <i>lower</i> than that of <i>ftp</i>.” (11) [SYM_P_0069290]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections.</p>	<p><u>103:</u></p> <p>“Like the host monitor, the LAN monitor consists of a <i>LAN event generator</i> (LEG) and a <i>LAN agent</i>. The LEG is currently a subset of UC Davis’ NSM [3]. Its main responsibility is to observe all of the traffic on its segment of the LAN to monitor host-to-host connections, services used, and volume of traffic. The LAN monitor reports on such network activity as <i>rlogin</i> and <i>telnet</i> connections, the use of security-related services, and changes in network traffic patterns.” (169) [SYM_P_0077177]</p> <p>“The host monitor consists of a <i>host event generator</i> (HEG) and a <i>host agent</i>. The HEG collects and analyzes audit records from the host’s operating system. The audit records are scanned for <i>notable events</i>, which are transactions that are of interest independent of any other records. These include, among others, failed events, user authentications, changes to the security state of the system, and any network access such as <i>rlogin</i> and <i>rsh</i>.” (169) [SYM_P_0077177]</p> <p>“The detection of certain attacks against a networked system of</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically.” (13) [SYM_P_0069292]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-79 to 4-80 [SYM_P_0075135-36].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>NFR. See Ranum et al., “Implementing a Generalized Tool for Network Monitoring,” Proceedings of the Eleventh Systems Administration Conference (LISA ’97), San Diego, CA, Oct. 1997 [SYM_P_0070720-28], 5-6 [SYM_P_0070725-26].</p>	<p>computers requires information from multiple sources. A simple example of such an attack is the so-called <i>doorknob</i> attack. In a doorknob attack the intruder’s goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers. These simple attacks can be remarkably successful [4]. As a case in point, UC Davis’ NSM recently observed an attacker of this type gaining super-user access to an external computer which did not require a password for the super-user account. In this case, the intruder used <i>telnet</i> to make the connection from a university computer system, and then repeatedly tried to gain access to several different computers at the external site. In cases like these, the intruder tries only a few logins on each machine (usually with different account names), which means that an IDS on each host may not flag the attack. Even if the behavior is recognized as an attack on the individual host, current IDS’s are generally unable to correlate reports from multiple hosts; thus they cannot recognize the <i>doorknob</i> attack as such. Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host’s monitor.” (168) [SYM_P_0077176]</p> <p>“In another incident, our NSM recently observed an intruder</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		SNMP/RMON. See my expert report.	<p>gaining access to a computer using a guest account which did not require a password. Once the attacker had access to the system, he exhibited behavior which would have alerted most existing IDS's (e.g., changing passwords and failed events). In an incident such as this, DIDS would not only report the attack, but may also be able to identify the source of the attack. That is, while most IDS's would report the occurrence of an incident involving user "guest" on the target machine, DIDS would also report that user "guest" was really, for example, user "smith" on the source machine, assuming that the source machine was in the monitored domain. It may also be possible to go even further back and identify all of the different user accounts in the "chain" to find the initial launching point of the attack." (168) [SYM_P_0077176]</p> <p>"The LAN monitor is currently a subset of UC Davis' Network Security Monitor [3]. The LAN monitor builds its own 'LAN audit trail'. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several simple analysis techniques to identify significant events. The</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
			<p>events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-79 to 4-80 [SYM_P_0075135-36].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>NFR. See Ranum et al., “Implementing a Generalized Tool for Network Monitoring,” Proceedings of the Eleventh Systems Administration Conference (LISA ’97), San Diego, CA, Oct. 1997 [SYM_P_0070720-28], 5-6 [SYM_P_0070725-26].</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
3	The method of claim 1, wherein the measure monitors data transfers by monitoring network packet data transfer errors.	<p>103:</p> <p>“The knowledge about capabilities of each of the network services is the third type of input (e-g., <i>telnet</i> provides the user with more capability than <i>ftp</i> does).” (11) [SYM_P_0069290]</p> <p>“The security level of the service is based on the capabilities of that service and the authentication required by that service. The <i>ftp</i> service, for example, has great capabilities with no authentication, so the security level for <i>ftp</i> is high. The <i>telnet</i> service, on the other hand, also has great capabilities, but it also requires strong authentication. Therefore, the security level for <i>telnet</i> is <i>lower</i> than that of <i>ftp</i>.” (11) [SYM_P_0069290]</p> <p>“One means of detecting anomalous behavior is to monitor statistical measures of user activities on the system. A popular way to monitor statistical measures is to keep <i>profiles</i> of legitimate user activities [2,6]. These profiles may include such items as login times, CPU usage, favorite editor and compiler, disk usage, number of printed pages per session, session length, error rate, etc. (T.F. Lunt et al., [7] presents a comprehensive list</p>	<p>SNMP/RMON. See my expert report.</p> <p>103:</p> <p>“The detection of certain attacks against a networked system of computers requires information from multiple sources. A simple example of such an attack is the so-called <i>doorknob</i> attack. In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers. These simple attacks can be remarkably successful [4]. As a case in point, UC Davis' NSM recently observed an attacker of this type gaining super-user access to an external computer which did not require a password for the super-user account. In this case, the intruder used <i>telnet</i> to make the connection from a university computer system, and then repeatedly tried to gain access to several different computers at the external site. In cases like these, the intruder tries only a few logins on each machine (usually with different account names), which means that an IDS on each host may not flag the attack. Even if the behavior is recognized as an attack on the individual host, current IDS's are generally unable to correlate reports from multiple hosts; thus they cannot</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>of possible measures.) The IDS will then use these profiles to compare current user activity with past user activity. Whenever a current user's activity pattern falls outside certain pre-defined thresholds, the behavior is considered anomalous. Legitimate behavior that is flagged as intrusive is defined to be a <i>false alarm</i>. A major problem with the statistical method is determining exactly what activities and statistical measures provide the highest detection rate and lowest false alarm rate for a particular system. Those statistics that detect an attack on a computer system may differ from system to system depending on the system and its environment; so the measures must be tailored for each particular system. It may also be the case that a particular activity may not be threatening by itself, but when aggregated with other activities, it may constitute an attack. These statistical profiles must be adaptive, i.e., they must be updated regularly, since users may be constantly changing their behavior.” (2) [SYM_P_0069281]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based</p>	<p>recognize the <i>doorknob</i> attack as such. Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host’s monitor.” (168) [SYM_P_0077176]</p> <p>“In another incident, our NSM recently observed an intruder gaining access to a computer using a guest account which did not require a password. Once the attacker had access to the system, he exhibited behavior which would have alerted most existing IDS’s (e.g., changing passwords and failed events). In an incident such as this, DIDS would not only report the attack, but may also be able to identify the source of the attack. That is, while most IDS’s would report the occurrence of an incident involving user “guest” on the target machine, DIDS would also report that user “guest” was really, for example, user “smith” on the source machine, assuming that the source machine was in the monitored domain. It may also be possible to go even further back and identify all of the different user accounts in the “chain” to find the initial launching point of the attack.” (168) [SYM_P_0077176]</p> <p>“The host monitor is currently installed on Sun SPARCstations running SunOS 4.0.x with the Sun C2 security package [9]. Through the C2 security package, the operating system produces</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically.” (13) [SYM_P_0069292]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-79 - 4-80 [SYM_P_0075135-36], 4-61 [SYM_P_0075117], 4-67 [SYM_P_0075123], 4-69 [SYM_P_0075125], 4-82 [SYM_P_0075138].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>SNMP/RMON. See my expert report.</p>	<p>audit records for virtually every transaction on the system. These transactions include file accesses, system calls, process executions, and logins. The contents of the Sun C2 audit record are: record type, record event, time, real user ID, audit user ID, effective user ID, real group ID, process ID, error code, return value, and label.” (170) [SYM_P_0077178]</p> <p>“The host monitor consists of a <i>host event generator</i> (HEG) and a <i>host agent</i>. The HEG collects and analyzes audit records from the host’s operating system. The audit records are scanned for <i>notable events</i>, which are transactions that are of interest independent of any other records. These include, among others, failed events, user authentications, changes to the security state of the system, and any network access such as <i>rlogin</i> and <i>rsh</i>.” (169) [SYM_P_0077177]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
			<p>simple analysis techniques to identify significant events. The events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-79 - 4-80 [SYM_P_0075135-36], 4-61 [SYM_P_0075117], 4-67 [SYM_P_0075123], 4-69 [SYM_P_0075125], 4-82 [SYM_P_0075138].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>SNMP/RMON. See my expert report.</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
4	The method of claim 1, wherein the measure monitors data transfers by monitoring network packet data transfer volume.	<p>“The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic.” (13) [SYM_P_0069292]</p> <p>“The abnormality of a connection is based on the probability of that particular connection occurring and the behavior of the connection itself. If a connection from host A to host B by service C is rare, then the abnormality of that connection is high. Furthermore, if the profile of that connection compared to a typical connection by the same type of service is unusual (e.g., the number of packets or bytes is unusually high for a <i>mail</i> connection), the abnormality of that connection is high.” (11) [SYM_P_0069290]</p> <p>“Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>.” (13) [SYM_P_0069292]</p>	<p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.” (171) [SYM_P_0077179]</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
5	The method of claim 1, wherein the measure monitors network connections by monitoring network connection requests.	<p>“The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic.” (13) [SYM_P_0069292]</p> <p>“The NSM models the network and hosts being monitored in the hierarchically-structured Interconnected Computing Environment Model (ICEM). The ICEM is composed of six layers, the lowest being the bit streams on the network, and the highest being a representation for the state of the entire networked system.</p> <p>The bottom-most, or first, layer is the <i>packet layer</i>. This layer accepts as input a <i>bit stream</i> from a broadcast local area network, viz. an Ethernet. The bit stream is divided up into complete Ethernet packets, and a time stamp is attached to the packet. This <i>time-augmented packet</i> is then passed up to the second layer.</p> <p>The next layer, called the <i>thread layer</i>, accepts as input the time-augmented packets from the packet layer. These packets are then correlated into unidirectional data streams. Each stream consists of the data (with the different layers of protocol headers removed) being transferred from one host to another host by a particular protocol (TCP/IP or UDP/IP), through a unique set</p>	<p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several simple analysis techniques to identify significant events. The events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>(for the particular set of hosts and protocol) of ports. This stream of data, which is called a thread, is mapped into a <i>thread vector</i>. All the thread vectors are passed up to the third layer.</p> <p>The <i>connection layer</i>, which is the third layer, accepts as input the thread vectors generated by the thread layer. Each thread vector is paired, if possible, to another thread vector to represent a bidirectional stream of data (i.e., a host-to-host connection). These pairs of thread vectors are represented by a connection vector generated by the combination of the individual thread vectors. Each connection vector will be analyzed, and a <i>reduced representation</i>, a <i>reduced connection vector</i>, is passed up to the fourth layer.</p> <p>Layer 4 is the <i>host layer</i> which accepts as input the reduced connection vectors generated by the connection layer. The connection vectors are used to build <i>host vectors</i>. Each host vector represents the network activities of a single host. These host vectors are passed up to the fifth layer.</p> <p>The <i>connected network</i> layer is the next layer in the ICEM hierarchy. It accepts as input the host vectors generated by the host layer. The host vectors are transformed into a graph G by treating the Data_path_tuples of the host vectors as an adjacency list. If G (host1,host2,serv1) is not empty, then there is a connection, or path, from host1 to host2 by service serv1. The</p>	

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>value for location $G(\text{host1}, \text{host2}, \text{serv1})$ is non empty if the host vector for host1 has $(\text{host2}, \text{serv1})$ in its Data_path_tuples. This layer can build the connected sub-graphs of G, called a <i>connected network vector</i>, and compare these sub-graphs against historical connected sub-graphs. This layer can also accept questions from the user about the graph. For example, the user may ask if there is some path between two hosts through any number of intermediate hosts — by a specific service. This set of connected network vectors is passed up to the sixth and final layer.</p> <p>The top most layer, called the <i>system layer</i>, accepts as input the set of connected network vectors from the connected network layer. The set of connected network vectors are used to build a <i>single system vector</i> representing the behavior of the entire system.” (10) [SYM_P_0069289]</p> <p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically. Like the host monitor, the LAN monitor provides an agent for communications with the director. In addition to handling queries of the audit data from the director, this agent gives the director access to a number of network management tools, which are analogous to the host operating system services provided by the host monitor.” (13) [SYM_P_0069292]</p>	

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
6	The method of claim 1, wherein the measure monitors network connections by monitoring network connection denials.	<p><u>103:</u></p> <p>“The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic.” (13) [SYM_P_0069292]</p> <p>“The NSM models the network and hosts being monitored in the hierarchically-structured interconnected Computing Environment Model (ICEM). The ICEM is composed of six layers, the lowest being the bit streams on the network, and the highest being a representation for the state of the entire networked system.</p> <p>The bottom-most, or first, layer is the <i>packet layer</i>. This layer accepts as input a <i>bit stream</i> from a broadcast local area network, viz. an Ethernet. The bit stream is divided up into complete Ethernet packets, and a time stamp is attached to the packet. This <i>time-augmented packet</i> is then passed up to the second layer.</p> <p>The next layer, called the <i>thread layer</i>, accepts as input the time-augmented packets from the packet layer. These packets are then correlated into unidirectional data streams. Each stream consists of the data (with the different layers of protocol headers</p>	<p><u>103:</u></p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several simple analysis techniques to identify significant events. The events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-72 [SYM_P_0075128], 4-62 [SYM_P_0075118].</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>removed) being transferred from one host to another host by a particular protocol (TCP/IP or UDP/IP), through a unique set (for the particular set of hosts and protocol) of ports. This stream of data, which is called a thread, is mapped into a <i>thread vector</i>. All the thread vectors are passed up to the third layer.</p> <p>The <i>connection layer</i>, which is the third layer, accepts as input the thread vectors generated by the thread layer. Each thread vector is paired, if possible, to another thread vector to represent a bidirectional stream of data (i.e., a host-to-host connection). These pairs of thread vectors are represented by a connection vector generated by the combination of the individual thread vectors. Each connection vector will be analyzed, and a <i>reduced</i> representation, a <i>reduced connection vector</i>, is passed up to the fourth layer.</p> <p>Layer 4 is the <i>host layer</i> which accepts as input the reduced connection vectors generated by the connection layer. The connection vectors are used to build <i>host vectors</i>. Each host vector represents the network activities of a single host. These host vectors are passed up to the fifth layer.</p> <p>The <i>connected network layer</i> is the next layer in the ICEM hierarchy. It accepts as input the host vectors generated by the host layer. The host vectors are transformed into a graph G by treating the Data path tuples of the host vectors as an adjacency</p>	<p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A [SYM_P_0078067-68].</p> <p>SNMP/RMON. See my expert report.</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>list. If $G(\text{host1}, \text{host2}, \text{serv1})$ is not empty, then there is a connection, or path, from host1 to host2 by service serv1. The value for location $G(\text{host1}, \text{host2}, \text{serv1})$ is non empty if the host vector for host1 has $(\text{host2}, \text{serv1})$ in its Data_path_tuples. This layer can build the connected sub-graphs of G, called a <i>connected network vector</i>, and compare these sub-graphs against historical connected sub-graphs. This layer can also accept questions from the user about the graph. For example, the user may ask if there is some path between two hosts through any number of intermediate hosts — by a specific service. This set of connected network vectors is passed up to the sixth and final layer.</p> <p>The top most layer, called the <i>system layer</i>, accepts as input the set of connected network vectors from the connected network layer. The set of connected network vectors are used to build a <i>single system</i> vector representing the behavior of the entire system.” (10) [SYM_P_0069289]</p> <p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically. Like the host monitor, the LAN monitor provides an agent for communications with the director. In addition to handling queries of the audit data from the director, this agent gives the director access to a number of network management tools, which are analogous to the host operating system services provided by the host monitor.” (13) [SYM_P_0069292].</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-72 [SYM_P_0075128], 4-62 [SYM_P_0075118].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A [SYM_P_0078067-68].</p> <p>SNMP/RMON. See my expert report.</p>	
7	The method of claim 1, wherein the measure monitors network connections by monitoring a correlation of network connections requests and network connection denials.	<p>103:</p> <p>“The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic.” (13) [SYM_P_0069292]</p> <p>“The NSM models the network and hosts being monitored in the hierarchically-structured Interconnected Computing Environment Model (ICEM). The ICEM is composed of six layers, the lowest being the bit streams on the network, and the</p>	<p>103:</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>highest being a representation for the state of the entire networked system.</p> <p>The bottom-most, or first, layer is the <i>packet layer</i>. This layer accepts as input a <i>bit stream</i> from a broadcast local area network, viz. an Ethernet. The bit stream is divided up into complete Ethernet packets, and a time stamp is attached to the packet. This <i>time-augmented packet</i> is then passed up to the second layer.</p> <p>The next layer, called the <i>thread layer</i>, accepts as input the time-augmented packets from the packet layer. These packets are then correlated into unidirectional data streams. Each stream consists of the data (with the different layers of protocol headers removed) being transferred from one host to another host by a particular protocol (TCP/IP or UDP/IP), through a unique set (for the particular set of hosts and protocol) of ports. This stream of data, which is called a thread, is mapped into a <i>thread vector</i>. All the thread vectors are passed up to the third layer.</p> <p>The <i>connection layer</i>, which is the third layer, accepts as input the thread vectors generated by the thread layer. Each thread vector is paired, if possible, to another thread vector to represent a bidirectional stream of data (i.e., a host-to-host connection). These pairs of thread vectors are represented by a connection vector generated by the combination of the individual thread</p>	<p>simple analysis techniques to identify significant events. The events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282].</p> <p>[SYM_P_0074948-5282], 1-10 [SYM_P_0074983], 4-63 [SYM_P_0075119], C-4 to C-5 [SYM_P_0075215-16], 4-62 [SYM_P_0075118], 4-72 [SYM_P_0075128].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>vectors. Each connection vector will be analyzed, and a <i>reduced</i> representation, a <i>reduced connection vector</i>, is passed up to the fourth layer.</p> <p>Layer 4 is the <i>host layer</i> which accepts as input the reduced connection vectors generated by the connection layer. The connection vectors are used to build <i>host vectors</i>. Each host vector represents the network activities of a single host. These host vectors are passed up to the fifth layer.</p> <p>The <i>connected network</i> layer is the next layer in the ICEM hierarchy. It accepts as input the host vectors generated by the host layer. The host vectors are transformed into a graph G by treating the Data_path_tuples of the host vectors as an adjacency list. If G (host1,host2,serv1) is not empty, then there is a connection, or path, from host1 to host2 by service serv1. The value for location G(host1,host2,serv1) is non empty if the host vector for host1 has (host2,serv1) in its Data_path_tuples. This layer can build the connected sub-graphs of G, called a <i>connected network vector</i>, and compare these sub-graphs against historical connected sub-graphs. This layer can also accept questions from the user about the graph. For example, the user may ask if there is some path between two hosts through any number of intermediate hosts — by a specific service. This set of connected network vectors is passed up to the sixth and final layer.</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
338		<p>The top most layer, called the <i>system layer</i>, accepts as input the set of connected network vectors from the connected network layer. The set of connected network vectors are used to build a <i>single system</i> vector representing the behavior of the entire system.” (10) [SYM_P_0069289]</p> <p>“Our previous work concentrated on the development of an intrusion-detection model and a prototype implementation of a network security monitor (NSM) for a broadcast local area network environment [5]. The NSM (adaptively) develops profiles of usage of network resources and then compares current usage patterns with the historical profile to determine possible security violations. The goal of our proposed research is to extend our network intrusion-detection concept from the LAN environment to arbitrarily wider areas with the network topology being arbitrary as well.” (11) [SYM_P_0069290]</p> <p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically. Like the host monitor, the LAN monitor provides an agent for communications with the director. In addition to handling queries of the audit data from the director, this agent gives the director access to a number of network management tools, which are analogous to the host operating system services provided by the host monitor.” (13) [SYM_P_0069292]</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282]. [SYM_P_0074948-5282], 1-10 [SYM_P_0074983], 4-63 [SYM_P_0075119], C-4 to C-5 [SYM_P_0075215-16], 4-62 [SYM_P_0075118], 4-72 [SYM_P_0075128].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p>	

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
8	The method of claim 1, wherein the measure monitors errors by monitoring error codes included in a network packet.	<p><u>103:</u></p> <p>“One means of detecting anomalous behavior is to monitor statistical measures of user activities on the system. A popular way to monitor statistical measures is to keep <i>profiles</i> of legitimate user activities [2,6]. These profiles may include such items as login times, CPU usage, favorite editor and compiler, disk usage, number of printed pages per session, session length, error rate, etc. (T.F. Lunt et al., [7] presents a comprehensive list of possible measures.) The IDS will then use these profiles to compare current user activity with past user activity. Whenever a current user's activity pattern falls outside certain pre-defined thresholds, the behavior is considered anomalous. Legitimate behavior that is flagged as intrusive is defined to be a <i>false alarm</i>. A major problem with the statistical method is determining exactly what activities and statistical measures provide the highest detection rate and lowest false alarm rate for a particular system. Those statistics that detect an attack on a computer system may differ from system to system depending on the system and its environment; so the measures must be tailored for each particular system. It may also be the case that a particular activity may not be threatening by itself, but when aggregated with other activities, it may constitute an attack. These statistical profiles must be adaptive, i.e., they must be updated regularly, since users may be constantly changing their behavior.” (2) [SYM_P_0069281]</p>	<p><u>103:</u></p> <p>“The detection of certain attacks against a networked system of computers requires information from multiple sources. A simple example of such an attack is the so-called <i>doorknob</i> attack. In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers. These simple attacks can be remarkably successful [4]. As a case in point, UC Davis' NSM recently observed an attacker of this type gaining super-user access to an external computer which did not require a password for the super-user account. In this case, the intruder used <i>telnet</i> to make the connection from a university computer system, and then repeatedly tried to gain access to several different computers at the external site. In cases like these, the intruder tries only a few logins on each machine (usually with different account names), which means that an IDS on each host may not flag the attack. Even if the behavior is recognized as an attack on the individual host, current IDS's are generally unable to correlate reports from multiple hosts; thus they cannot recognize the <i>doorknob</i> attack as such. Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host's monitor.”</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>“The DIDS LAN monitor is built on the same foundation as UC Davis' Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically.” (13) [SYM_P_0069292]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User's Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-67 [SYM_P_0075123], 4-82 [SYM_P_0075138].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and</p>	<p>(168) [SYM_P_0077176]</p> <p>“In another incident, our NSM recently observed an intruder gaining access to a computer using a guest account which did not require a password. Once the attacker had access to the system, he exhibited behavior which would have alerted most existing IDS's (e.g., changing passwords and failed events). In an incident such as this, DIDS would not only report the attack, but may also be able to identify the source of the attack. That is, while most IDS's would report the occurrence of an incident involving user “guest” on the target machine, DIDS would also report that user “guest” was really, for example, user “smith” on the source machine, assuming that the source machine was in the monitored domain. It may also be possible to go even further back and identify all of the different user accounts in the “chain” to find the initial launching point of the attack.” (168) [SYM_P_0077176]</p> <p>“The host monitor is currently installed on Sun SPARCstations running SunOS 4.0.x with the Sun C2 security package [9]. Through the C2 security package, the operating system produces audit records for virtually every transaction on the system. These transactions include file accesses, system calls, process executions, and logins. The contents of the Sun C2 audit record are: record type, record event, time, real user ID, audit user ID, effective user ID, real group ID, process ID, error code, return</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>SNNP/RMON. See my expert report.</p> <p>Jeremy Frank, “Artificial Intelligence and Intrusion Detection: Current and Future Directions,” Proc. of the 17th National Computer Security Conference (1994) [SYM_P_0073569-80].</p>	<p>value, and label.” (170) [SYM_P_0077178]</p> <p>“The host monitor consists of a <i>host event generator</i> (HEG) and a <i>host agent</i>. The HEG collects and analyzes audit records from the host’s operating system. The audit records are scanned for <i>notable events</i>, which are transactions that are of interest independent of any other records. These include, among others, failed events, user authentications, changes to the security state of the system, and any network access such as <i>rlogin</i> and <i>rsh</i>.” (169) [SYM_P_0077177]</p> <p>“The LAN monitor is currently a subset of UC Davis’ Network Security Monitor [3]. The LAN monitor builds its own ‘LAN audit trail’. The LAN monitor observes each and every packet on its segment of the LAN and, from these packets, it is able to construct higher-level objects such as connections (logical circuits), and service requests using the TCP/IP or UDP/IP protocols. In particular, it audits host-to-host connections, services used, and volume of traffic per connection.</p> <p>Similar to the host monitor, the LAN monitor uses several simple analysis techniques to identify significant events. The events include the use of certain services (e.g., <i>rlogin</i> and <i>telnet</i>) as well as activity by certain classes of hosts (e.g., a PC without a host monitor). The LAN monitor also uses and maintains profiles of expected network behavior. The profiles consist of</p>

**Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”**

Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
			<p>expected data paths (e.g., which systems are expected to establish communication paths to which other systems, and by which service) and service profiles (e.g., what a typical <i>telnet</i>, <i>mail</i>, or <i>finger</i> is expected to look like).” (171) [SYM_P_0077179]</p> <p>See L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. See NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-67 [SYM_P_0075123], 4-82 [SYM_P_0075138].</p> <p>ISS RealSecure. See Real Secure 1.1: User Guide and Reference Manual (March 1997) [SYM_P_0078004-77], Appendix A.</p> <p>SNNMP/RMON. See my expert report.</p> <p>Jeremy Frank, “Artificial Intelligence and Intrusion Detection: Current and Future Directions,” Proc. of the 17th National Computer Security Conference (1994) [SYM_P_0073569-80].</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

'338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
10	The method of claim 8 wherein an error code comprises an error code indicating a reason a packet was rejected.	<p><u>103:</u></p> <p>“One means of detecting anomalous behavior is to monitor statistical measures of user activities on the system. A popular way to monitor statistical measures is to keep <i>profiles</i> of legitimate user activities [2,6]. These profiles may include such items as login times, CPU usage, favorite editor and compiler, disk usage, number of printed pages per session, session length, error rate, etc. (T. F. Lunt et al., [7]) presents a comprehensive list of possible measures.) The IDS will then use these profiles to compare current user activity with past user activity. Whenever a current user's activity pattern falls outside certain pre-defined thresholds, the behavior is considered anomalous. Legitimate behavior that is flagged as intrusive is defined to be a <i>false alarm</i>. A major problem with the statistical method is determining exactly what activities and statistical measures provide the highest detection rate and lowest false alarm rate for a particular system. Those statistics that detect an attack on a computer system may differ from system to system depending on the system and its environment; so the measures must be tailored for each particular system. It may also be the case that a particular activity may not be threatening by itself, but when aggregated with other activities, it may constitute an attack. These statistical profiles must be adaptive, i.e., they must be updated regularly, since users may be constantly changing their behavior.” (2) [SYM P. 0069281]</p>	<p><u>103:</u></p> <p>“The detection of certain attacks against a networked system of computers requires information from multiple sources. A simple example of such an attack is the so-called <i>doorknob</i> attack. In a doorknob attack the intruder's goal is to discover, and gain access to, insufficiently-protected hosts on a system. The intruder generally tries a few common account and password combinations on each of a number of computers. These simple attacks can be remarkably successful [4]. As a case in point, UC Davis' NSM recently observed an attacker of this type gaining super-user access to an external computer which did not require a password for the super-user account. In this case, the intruder used <i>telnet</i> to make the connection from a university computer system, and then repeatedly tried to gain access to several different computers at the external site. In cases like these, the intruder tries only a few logins on each machine (usually with different account names), which means that an IDS on each host may not flag the attack. Even if the behavior is recognized as an attack on the individual host, current IDS's are generally unable to correlate reports from multiple hosts; thus they cannot recognize the <i>doorknob</i> attack as such. Because DIDS aggregates and correlates data from multiple hosts and the network, it is in a position to recognize the doorknob attack by detecting the pattern of repeated failed logins even though there may be too few on a single host to alert that host's monitor.”</p>

Distributed Intrusion Detection System
“DIDS February 1991 and DIDS October 1991”

‘338 Claim number	Claim Term	DIDS February 1991 (printed publication and public use)	DIDS October 1991 (printed publication and public use)
		<p>“The DIDS LAN monitor is built on the same foundation as UC Davis’ Network Security Monitor [5]. Since there is no native LAN audit trail, the LAN monitor is responsible for building its own. The LAN monitor sees every packet on its segment of the LAN. From these packets, the LAN monitor is able to construct higher level objects such as connections (logical circuits), and service requests. In particular, it audits host-to-host connections, services used, and volume of traffic. Like the host based monitor, the LAN monitor uses several levels of analysis to catch the most significant events, for example, sudden changes in network load, the use of security-related services, and network activities such as <i>rlogin</i>. As with the host monitor, the LAN monitor retains the audit data for analysis by the director. It also uses and maintains profiles of network behavior, which are updated periodically.” (13) [SYM_P_0069292]</p> <p><i>See</i> L. Todd Heberlein, “A Network Security Monitor – Final Report” (1995) [SYM_P_0070787-839], 51-53 [SYM_P_0070837-39] (public use).</p> <p>NetRanger. <i>See</i> NetRanger User’s Guide 1.3.1, WheelGroup Corporation, 1997 [SYM_P_0074948-5282], 4-67 [SYM_P_0075123].</p> <p>ISS RealSecure. <i>See</i> Real Secure 1.1: User Guide and</p>	<p>(168) [SYM_P_0077176]</p> <p>“In another incident, our NSM recently observed an intruder gaining access to a computer using a guest account which did not require a password. Once the attacker had access to the system, he exhibited behavior which would have alerted most existing IDS’s (e.g., changing passwords and failed events). In an incident such as this, DIDS would not only report the attack, but may also be able to identify the source of the attack. That is, while most IDS’s would report the occurrence of an incident involving user “guest” on the target machine, DIDS would also report that user “guest” was really, for example, user “smith” on the source machine, assuming that the source machine was in the monitored domain. It may also be possible to go even further back and identify all of the different user accounts in the “chain” to find the initial launching point of the attack.” (168) [SYM_P_0077176]</p> <p>“The host monitor is currently installed on Sun SPARCstations running SunOS 4.0.x with the Sun C2 security package [9]. Through the C2 security package, the operating system produces audit records for virtually every transaction on the system. These transactions include file accesses, system calls, process executions, and logins. The contents of the Sun C2 audit record are: record type, record event, time, real user ID, audit user ID, effective user ID, real group ID, process ID, error code, return</p>